

ALMA MATER STUDIORUM - UNIVERSITÀ DI BOLOGNA
CESENA CAMPUS

DIPARTIMENTO DI DEPARTMENT OF ELECTRICAL,
ELECTRONIC, AND INFORMATION ENGINEERING -
“Guglielmo Marconi” – DEI”

SECOND CYCLE DEGREE IN BIOMEDICAL ENGINEERING
Class: LM-21

THESIS TITLE

**A NEURAL MASS MODEL OF SELECTIVE ATTENTION OVER
SEMANTIC MEMORY**

Graduation thesis in
Neural systems

SUPERVISOR

PROF. MAURO URSINO

CANDIDATE

RAMEEN SHAIKH

Academic Year: 2025-2026

Session: 25/09/2026

TABLE OF CONTENT

ABSTRACT.....	1
INTRODUCTION.....	2
CHAPTER 1: LITERATURE REVIEW	4
1.1 BRAIN OSCILLATIONS AND COGNITION	4
1.1.1 The gamma rhythm.....	4
1.1.2 The alpha rhythm and inhibitory control	5
1.1.3 Alpha-gamma coupling	5
1.2 STATE OF THE ART	6
1.2.1 Neural mass framework.....	6
1.2.2 Models of alpha-mediated attention	7
1.2.3 Models of semantic memory	7
1.2.4 Short-term synaptic plasticity	8
1.2.5 The present contribution.....	9
1.3 PREVIOUS MODELS	9
CHAPTER 2: METHODOLOGY	15
2.1 THE OLD MODEL	15
2.1.1 The Neural Mass Model of a Single Cortical Column	15
2.1.2 The Animal Taxonomy.....	18
2.1.3 Network Training	22
2.2 NEW MODEL	24
2.2.1 Multi-region architecture	24
2.2.2 Inhibition Mechanism.....	26
2.2.3 Spotlight Mechanism.....	26
2.2.4 Fatigue Mechanism.....	26
CHAPTER 3: SIMULATIONS AND RESULTS	28
3.1 SENSITIVITY ANALYSIS AND PARAMETER VALIDATION	28
3.2 SELECTIVE ATTENTION – SPOTLIGHT MECHANISM.....	29
3.2.1 A fully correct case.....	31
3.2.2 A difficult concept: the giraffe (Giraffe–Cow–Pigeon)	34
3.2.3 Competing demands: the Giraffe–Zebra–Lion triplet	35
3.2.4 Recall from a marginal feature	38
3.2.5 Recovering the rooster window: Hen–Goose–Rooster	39
3.3 SUMMARY	42

CHAPTER 4: DISCUSSION	43
4.1 APPLYING ALPHA GATING TO A TRAINED SEMANTIC MEMORY	43
4.2 SEQUENTIAL RETRIEVAL AND THE ROLE OF FATIGUE.....	43
4.3 ROBUSTNESS AND THE STRUCTURE OF THE DIFFICULT CASES.....	44
4.4 LIMITATIONS.....	44
4.5 FUTURE WORK.....	45
4.6 SUMMARY	46
REFERENCES.....	47

ABSTRACT

Semantic memory and selective attention are usually studied apart. But they share the same oscillatory machinery in the brain, and this thesis brings them together in a single neural mass model. It extends a trained semantic-memory network with an alpha-gated attention mechanism, so stored concepts can be selected and retrieved one at a time.

The model has two starting points. One is a semantic-memory network, trained with an asymmetric Hebbian rule on 102 features across fourteen animals. The other is an alpha-rhythm gating mechanism, which suppresses information through phase opposition between regions. These sit inside a three-region architecture: a sensory area, devoted to feature representation; an alpha oscillator; and a semantic recall area, which contains the auto-associative synapses and the fatigue mechanism. On top of this, the thesis adds three mechanisms. An inhibition mechanism applies alpha suppression to the sensory network. A spotlight mechanism releases that suppression selectively, moving attention from one concept to the next. A fatigue mechanism, modelled as activity-dependent depletion of recurrent excitation, lets an active concept fade so the network can move on.

We simulated a situation in which three animals are presented in sequence, each cued by a single distinctive feature in the sensory area. The attentional spotlight mechanism correctly segments the three animals in the majority of cases: each one is retrieved individually, within its own 400ms temporal window, with all of its features correctly recovered in the semantic recall area. Only in a few complex cases, where an animal shared many features with others in the taxonomy, did the model need a small adjustment to retrieve everything correctly. Increasing the excitation driving the retrieval was usually enough to fix this. In one case, though, this same fix caused a new problem: raising the excitation enough to fully retrieve two animals with many shared features pushed a third, less overlapping animal into saturation, and increasing inhibition only partly solved it. Another case showed that an animal's retrieval could also depend on its position in the sequence rather than on its own features, which suggests that timing, not just shared content, can make retrieval difficult.

In conclusion, although the alpha gating mechanism was originally built for suppressing distractors, our simulations demonstrate that it can be used to select and retrieve stored knowledge; a simple fatigue mechanism is enough to keep that retrieval moving on its own, without the need for any additional external mechanism.

Keywords: Neural mass model, semantic memory, selective attention, alpha rhythm, synaptic fatigue, computational neuroscience.

INTRODUCTION

The human ability to recognize objects, retrieve their semantic content, and selectively focus on relevant information while suppressing irrelevant distractors represents one of the most fundamental aspects of cognition. The two processes – semantic memory retrieval and selective attention have traditionally been studied as separate systems. However, increasing experimental and theoretical evidence shows that they are intertwined and share a common neural substrate (Nobre & Stokes, 2019). This thesis presents a neurocomputational model that bridges the two processes. It extends an existing model of semantic memory with a mechanism for selective attention based on the gating of information by the alpha rhythm, so that stored concepts can be accessed one at a time under attentional control.

Semantic memory is the conceptual knowledge about objects and their characteristics that humans build up throughout their lives. In this model, the concepts are embedded within a hierarchical taxonomy, such that individual exemplars (e.g., lion, zebra) inherit properties from intermediate categories (e.g., wild, animal, mammal) and from superordinate ones (e.g., animal). Furthermore, not all features of a concept carry equal weight; some concepts are salient and reliably evoked, while others are marginal and only rarely recalled. Capturing this hierarchical structure in a biologically plausible neural network is the core modelling challenge, addressed in the foundational work by (Ursino & Pirazzini, 2024) which serves as a basis of this thesis.

In real cognitive scenarios, semantic retrieval unfolds against a background of competing representations that must be gated, ignored, or processed sequentially. Accessing stored semantic knowledge is not a passive look-up process but an active, attention-dependent operation that recruits the same neural pathways used to select and suppress perceptual representations. The alpha rhythm (8-12 Hz) has been proposed as a key mechanism for such attentional suppression. Importantly, not only the amplitude but also the phase of the alpha rhythm is essential. Taking that into account, (Ursino, 2025) proposed a neural mass model in which an alpha-generating region transmits an out-of-phase signal to sensory areas coding for distractors, effectively preventing their transmission to downstream detection areas.

The present thesis brings these two ideas together. It places a trained semantic memory inside a three-region architecture and uses alpha-based gating to control which stored concept is retrieved. The first region (ROI 1) is a semantic store of 102 neural mass units, each coding one feature of a concept. The second region (ROI 3) is a dedicated alpha oscillator that broadcasts a rhythmic inhibitory signal to the rest of the network. The third region (ROI 4) is an auto-associative recall area in which the trained recurrent connections reconstruct a full concept from a partial cue¹. Attention is implemented as a spotlight: the alpha rhythm

¹ The numbering skips ROI 2 to remain consistent with the foundational model of (Ursino, 2025), which included two sensory areas, ROI 1 and ROI 2. The present model retains only ROI 1.

suppresses the entire store by default, and a concept is selected by releasing that suppression from its cue feature alone, which allows the recurrent network to complete the concept while all other representations remain gated. Moving the released window from one feature to another shifts attention across the network in time.

In the model, each semantic feature is represented by a neural mass unit, built from interconnected populations of pyramidal neurons, excitatory interneurons, and slow and fast inhibitory interneurons arranged in feedback, and tuned to oscillate in the gamma band (30–40 Hz), consistent with the gamma activity observed during semantic memory tasks (David & Friston, 2003; Jansen & Rit, 1995). The network was trained on 14 animals — eight mammals (dog, cat, sheep, cow, giraffe, zebra, bear, lion) and six birds (owl, rooster, hen, pigeon, goose, parrot).

A further contribution of the thesis is the introduction of synaptic fatigue in the recall area, modelled as the short-term depletion of the available transmitter pool at the recurrent synapses. Sustained activity in the unit gradually depletes its resources and weakens the recurrent drive that sustains the active concept, so that a representation fades once it has been held for a few hundred milliseconds. Combined with the spotlight mechanism, this fatigue allows the network to disengage from one concept and take up the next without an external reset, producing a clean sequence of distinct retrievals from minimal input and without interference between them. Together, alpha gating and synaptic fatigue offer a biologically grounded account of how the brain might access stored knowledge serially, one concept at a time.

This thesis is organized as follows. Chapter 1 reviews the relevant background, covering brain oscillations and their role in cognition, the functional roles of the gamma and alpha rhythms and their coupling, and the neural mass modelling tradition from the Wilson–Cowan models through to the Jansen–Rit column on which the present model is built. Chapter 2 sets out the methodology in two parts: the existing model that this work starts from, comprising the single-column neural mass unit, the animal taxonomy, and the Hebbian training procedure; and the new model, comprising the three-region architecture together with the inhibition, spotlight, and fatigue mechanisms. Chapter 3 presents the simulations and results, beginning with a sensitivity analysis of the model's robustness to noise and parameter choices and continuing with the spotlight-based retrieval of concepts across a range of animal groupings. Chapter 4 discusses the findings, relates them to the existing literature, sets out the limitations of the model, and outlines directions for future work.

CHAPTER 1: LITERATURE REVIEW

1.1 BRAIN OSCILLATIONS AND COGNITION

Neurons across cortical and subcortical brain areas exhibit coordinated frequency bands, and each band corresponds to a different functional state. These rhythms shape how the brain handles information: they open temporal windows of excitability and inhibition, and coordinate the communication between regions that are anatomically far (Buzsaki & Draguhn, 2004; Fries, 2015). The bands are conventionally divided into five: Delta (0.5–4 Hz), theta (4–7 Hz), Alpha (8–13 Hz), Beta (13–30 Hz), Gamma (>30 Hz).

Two of these bands are central to the present work. The gamma rhythm (here taken as the slow gamma range, ~30–40 Hz) is associated with the active representation of information in local cortical circuits and with the binding of distributed features into coherent objects. The alpha rhythm (~8–13 Hz) is associated with top-down inhibitory control and the suppression of task-irrelevant populations. Understanding the functional roles and interactions of these two rhythms is essential for interpreting the model developed in this thesis.

1.1.1 The gamma rhythm

Slow gamma oscillations between 30 and 40 Hz play a significant role in attention, arousal, object recognition, and language perception tasks (Jensen & Colgin, 2007; Tallon-Baudry & Bertrand, 1999). Its most influential interpretation concerns the binding of features. Gray and Singer showed that spatially separated neurons responding to features of the same visual object tend to fire synchronously in the gamma band (Gray, 1989; Singer & Gray, 1995). This finding is interpreted as a mechanism of feature binding, by which the distributed representations of an object's attributes are integrated into a single percept.

Gamma activity also relates to memory. (Tallon-Baudry & Bertrand, 1999) reported stronger gamma-band responses when participants recognized familiar objects than when they viewed novel objects. Familiar objects already have representations in long-term memory. Herrmann and colleagues developed this idea into the match-and-utilization model. According to this model, gamma activity increases when incoming sensory information matches a representation already stored in memory (Herrmann, 2004). Gamma responses have also been observed in temporal and prefrontal regions during object naming, category judgements, and conceptual priming.

These observations motivate the representational choice made in this thesis. Following (Ursino et al., 2010; Ursino & Pirazzini, 2024), each unit that codes a semantic feature is implemented as a four-population cortical column whose parameters place it in the gamma range, so that the presence of a feature in a concept is signaled by gamma-band activity, and the co-activation of features belonging to the same concept appears as gamma synchrony.

1.1.2 The alpha rhythm and inhibitory control

The alpha rhythm is defined as oscillatory activity in the 8-13 Hz range. It was the first brain oscillation described systematically by Hans Berger in 1929. For much of the following century, alpha was interpreted as a sign of cortical idling, a default state of reduced activity seen when subjects closed their eyes or sat idle (Pfurtscheller, 1996). However, Evidence accumulated over the past two decades has substantially revised this view. Recordings from EEG, MEG, and intracranial electrodes now support a more active and functionally specific role for alpha, i.e., it actively inhibits neural populations that are not relevant to the current task (Bacigalupo & Luck, 2022; Janssens, 2018; van Diepen, 2019; Wildegger, 2017; Wostmann, 2019).

Klimesch formalized this inhibitory role, proposing that alpha implements a pulsed inhibition that rhythmically modulates cortical excitability (Klimesch, 2012). On this account, the troughs of the alpha cycle are brief windows during which stimulus-driven activity can propagate, while the peaks impose strong inhibition that blocks processing. By controlling the phase and amplitude of alpha in particular regions, the brain can place an attentional gate that is precise in both space and time. The experimental support is substantial. EEG and MEG studies report that alpha power decreases (event-related desynchronization) over regions contralateral to an attended stimulus and increases (event-related synchronization) over regions ipsilateral to the attended stimulus and contralateral to distractors (Ikkai, 2016; Thut, 2006; Worden, 2000; Wostmann, 2019). Crucially, both the phase and the amplitude of the rhythms are important. Therefore, the phase of alpha at which a stimulus arrives predicts whether it is detected or not (Busch, 2009; Mathewson, 2009).

A key point, developed by (Ursino, 2025) and central to the present model, is that alpha must reach sensory areas in phase opposition to the signal transmitted to the detection area to achieve effective distractor suppression. When the two signals are roughly 180 degrees out of phase, the sensory area is maximally active precisely when the detection area is maximally inhibited, so its activity cannot drive downstream processing. In the present model, the alpha rhythm is produced by a dedicated neural mass unit whose anatomical identity is left unspecified; what matters functionally is that it generates a sustained ~ 10 Hz rhythm and transmits it to downstream regions with a controlled phase.

1.1.3 Alpha-gamma coupling

Over the past decade, cross-frequency coupling between oscillatory bands has been recognized as a mechanism behind higher-order cognition. The coupling between alpha and gamma is of particular relevance here. The amplitude of gamma activity is modulated by the phase of the ongoing alpha rhythm (Jensen & Colgin, 2007; Voytek, 2010) which provides a natural substrate for alpha to gate gamma-band processing.

(Fries, 2015) framed this in information terms with the communication-through-coherence hypothesis: two populations communicate well when they oscillate in synchrony, with a stable phase relationship, because excitatory inputs that arrive at the optimal phase of the local oscillation have the largest effect on the target. Desynchronization, or an imposed phase shift, weakens transmission between the populations. The alpha-gating mechanism used in this thesis follows directly from this idea. By imposing a phase opposition between the alpha in the sensory area and the alpha in the detection area, the model decouples their gamma-band communication and suppresses the transmission of distractor representations.

(Ursino, 2025) implemented this idea in a neural mass network, showing that an alpha-suppression mechanism based on phase opposition could realize selective attention, and a sensitivity analysis identified the phase difference between the two alpha signals as the parameter most critical to successful suppression. This phase-opposition mechanism is the foundation on which the attentional gating in the present model is built.

1.2 STATE OF THE ART

The thesis presents a computational model that bridges two areas of research that are usually kept separate: semantic memory and selective attention. Both have been studied with neural mass models, the framework adopted here, but rarely in the same network. This section reviews the state of the art in each area, the modelling framework they share, and identifies the gap that the present work fills.

1.2.1 Neural mass framework

Computational models of brain activity span a wide range of descriptions, and the right level depends on the question. Detailed biophysical models, following (Hodgkin & Huxley, 1952), reproduce the membrane dynamics of single neurons accurately but become intractable for a population of many thousands of cells. Abstract rate and field models sit at the opposite extreme, describing population activity economically but discarding the synaptic and cellular detail that gives the dynamics their physiological meaning. Neural mass models fall between the two. They describe the mean activity of a few interacting populations within a cortical column, which is efficient enough to simulate large networks, while keeping enough biophysical structure that their parameters and outputs can be related to real synaptic properties and to EEG and MEG recordings (David & Friston, 2003; Jansen & Rit, 1995). Most current modelling of cortical oscillations in cognition is done at this level, and it is the framework used here for both parts of the thesis.

Work in this framework has shown how different rhythms arise from the interaction of populations with different synaptic time constants. Gamma activity in particular depends on fast GABAergic inhibition, through the interaction of pyramidal cells and fast interneurons (the

PING mechanism) and the self-inhibition of the fast interneuron population (the ING mechanism) (Ursino et al., 2010; Wendling, 2002). Building on this, (Ursino et al., 2010) developed a four-population cortical column, with pyramidal neurons, excitatory interneurons, and slow and fast inhibitory interneurons, in which a fast-inhibitory self-loop produces robust oscillation in the gamma range. This column is the computational unit used throughout the present model.

1.2.2 Models of alpha-mediated attention

The first of the two domains concerns the role of the alpha rhythm in attention. Alpha was traditionally viewed as a sign of cortical idling. More recent research suggests that it actively controls cortical processing, and recent models aim to explain how this mechanism works. Several models stress that it is the phase of the alpha rhythm, and not only its amplitude, that controls the transmission of information. Jensen and colleagues proposed that the duty cycle of the alpha rhythm sets how many representations can be processed within a cycle, with the most salient appearing at the earliest phase (Jensen, 2012, 2014). (Bonnefond, 2017; Quax, 2017) argued that communication between regions is blocked when their alpha oscillations are asynchronous, linking alpha phase directly to the efficacy of gamma-band transmission, in line with the communication-through-coherence framework (Fries, 2015).

(Ursino, 2025) tested this principle using a neural mass network for selective attention. The model showed that the phase relationship between regions was critical for suppressing distractors. The mechanism was demonstrated in both a small illustrative network and a larger spatial-attention task with Gestalt-organized representations. However, the stimulus patterns and connectivity were externally specified rather than learned, and the resulting representations carried no learned semantic content. The present thesis addresses this limitation by applying attentional gating to a trained semantic-memory network.

1.2.3 Models of semantic memory

For the computational modelling of semantic memory, two broad traditions can be distinguished. The first uses feedforward networks trained by backpropagation to study how distributed conceptual representations emerge in hidden layers, beginning with the work of Rumelhart and Rogers and McClelland and continuing in the distributed-plus-hub models of Rogers and colleagues, which reconcile category-specific and category-general accounts of semantic organization (Chen, 2017; Rogers & McClelland, 2008). The second tradition uses attractor networks trained with Hebbian learning, in which a concept is stored as a stable state of a recurrent network; this line includes the feature-based models of McRae and colleagues and the multilayer, neuroanatomically grounded models of Garagnani and Pulvermüller, some versions of which use spiking neurons to study gamma-band activity during word processing (Garagnani, 2017; Garagnani & Pulvermüller, 2016; McRae, 1997).

A difficulty common to much of this work is capturing the hierarchical organization of concepts and the unequal importance of features, since classical Hebbian rules tend to produce symmetric connections that do not distinguish superordinate from subordinate concepts. (Ursino & Pirazzini, 2024) addressed this with an attractor network of neural mass units in which each unit codes one feature and oscillates in the gamma band, trained with an asymmetric, presynaptically gated Hebbian rule. After training on a fourteen-animal taxonomy, the network forms categories, distinguishes subordinate from superordinate concepts, and separates salient from marginal features according to how often they occur during learning. This model provides the semantic store on which the present thesis builds.

1.2.4 Short-term synaptic plasticity

A third component of the present model, distinct from the two domains above, draws on the literature on short-term synaptic plasticity. This mechanism is distinguished from the Hebbian learning rule described in the previous section. Hebbian learning produces long-term changes in the synaptic weights during training and allows semantic information to be stored. Synaptic fatigue, in contrast, produces a temporary reduction in synaptic efficacy during recall and does not modify the learned connections permanently.

The depressing form of short-term synaptic plasticity, known as short-term synaptic depression, arises mainly from the depletion of the readily releasable pool of neurotransmitter vesicles. This pool is consumed by sustained activity and gradually recovers once activity falls (Abbott, 1997; Tsodyks & Markram, 1997). The standard phenomenological description, due to Tsodyks and Markram, represents this process with a resource variable that is depleted in proportion to presynaptic activity and recovers toward its baseline with a fixed time constant.

Functionally, short-term depression acts as a form of gain control and adaptation: by weakening the output of a population that has been strongly active, it prevents that population from dominating indefinitely (Abbott, 1997). This property has recently been exploited to drive autonomous transitions between stored states in memory models. In the consolidation model of (Singh, 2022), short-term synaptic depression gradually destabilizes an active attractor, ending its dominance so that the network can move to another memory without any external input. The fatigue mechanism in the present model uses the same mechanism for a different purpose: a depletion variable weakens the recurrent drive that sustains an active concept in the recall area, so that a retrieved representation fades after it has been held briefly, and the network becomes free to retrieve the next concept under the moving attentional spotlight. The mechanism is therefore not introduced ad hoc but applies a well-characterized property of cortical synapses to the problem of sequential retrieval.

1.2.5 The present contribution

Within the neural mass framework, attention models account for the selection and suppression of information but operate on networks without learned content, while semantic-memory models account for the structure of conceptual knowledge but do not address how attention selects among stored concepts. Both (Ursino & Pirazzini, 2024) and (Ursino, 2025) identified the integration of these two capabilities as a direction for future work. The present thesis carries out that integration. It places a trained semantic memory within a multi-region architecture and applies alpha-based phase-opposition gating to it, so that a moving attentional spotlight retrieves stored concepts one at a time, and a synaptic fatigue mechanism allows the network to move between them. In doing so, it joins the selection of information, studied in the attention literature, with the storage and structure of information, studied in the semantic-memory literature, within a single oscillatory framework.

1.3 PREVIOUS MODELS

The most detailed models of neural activity work at the level of the single neuron. The best-known is the Hodgkin–Huxley model (Hodgkin & Huxley, 1952), which describes a neuron's membrane potential through the voltage-dependent flow of sodium, potassium, and leak currents and reproduces the action potential in fine detail. Later single-neuron models simplify this picture to different degrees: the integrate-and-fire model reduces the neuron to a threshold device, the FitzHugh–Nagumo model to a two-variable approximation of its dynamics (FitzHugh, 1961; Nagumo, 1962), and the spiking-rate model replaces the precise timing of action potentials with an average firing rate passed through a sigmoidal nonlinearity. What these models share is that each still represents a neuron as an individual unit.

Neural mass models start from a simple observation about how the cortex is built. Neurons are not arranged at random but organized into vertical columns, functional units roughly 0.5–1 mm wide that run through all six cortical layers and contain on the order of tens of thousands of cells. Within a column, neurons of the same type tend to receive similar inputs and to fire with similar statistics, so rather than track each one separately, the whole population can be treated as a single unit and described by two average quantities: its mean membrane potential and its mean firing rate.

These two quantities are linked by two transformations. The mean firing rate of a population follows from its mean membrane potential through a sigmoidal function, which saturates at low and high potentials and is steepest in between, so that a population is silent when weakly driven, responds gradually as its input grows, and approaches a ceiling when strongly driven. The influence of one population on another runs in the opposite direction: the presynaptic firing rate is passed through a synaptic response function that reproduces the shape and time course of the postsynaptic potential, turning a rate back into a potential. A neural mass unit is built by

connecting a few such populations in feedback, and the interplay between excitation and inhibition across these connections is what produces oscillatory activity.

The appeal of this approach is that it captures the dynamics of a cortical area at a cost low enough to simulate large networks. A single unit reproduces the rhythmic activity of a population in a form that can be compared with local field potentials or EEG. Because each unit is inexpensive, hundreds of them can be connected to model a process that spans many cortical areas, as in the present work, where each unit stands for one semantic feature. Population models of this kind have been developed since the 1970s, and the following sections discuss two population-level models relevant to this thesis: the Wilson–Cowan model and the Jansen–Rit column.

1.3.1.a Wilson-Cowan model

The Wilson-Cowan model (Wilson & Cowan, 1972) is a mathematical description of the interaction between excitatory and inhibitory populations. Its goal is to simulate the macroscopic behaviour of neuronal networks with no attention to detail on singular neuronal contributions. It was the first model to treat excitatory–inhibitory feedback as the basic unit of population dynamics, the principle on which the later models in this section are built.




Figure 1.1 Pictorial depiction of the Wilson-Cowan model

The above figure depicts the Wilson-Cowan model. Here, the activity of excitatory populations is described by variable E , whereas the activity of inhibitory populations is represented by I . Each population sends a synapse both to itself and to the other population. The corresponding differential equations are:

$$\tau_e \frac{dE(t)}{dt} = -E(t) + S_e(aE(t) - bI(t) + i_e(t)) \quad (1.1)$$

$$\tau_i \frac{dI(t)}{dt} = -I(t) + S_i(cE(t) - dI(t) + i_i(t)) \quad (1.2)$$

Where τ_e and τ_i denote the time constants of the two populations, and a , b , c , and d are the four synaptic weights of the scheme above. The terms i_e and i_i are the external inputs, taken here as constant. The functions S_e and S_i are sigmoidal activation functions returning the fraction of active neurons, with 0 corresponding to complete inhibition and 1 to complete excitation:

$$S_j(u) = \frac{1}{1 + e^{-Y_j(u - u_{0j})}} \quad j \in \{e, i\} \quad (1.3)$$

Where Y_j is proportional to the central gain, while u_{0j} is the abscissa at the central point of sigmoid for population j . For sufficiently high input current and carefully chosen synaptic weights, the system exhibits an asymptotically stable limit cycle, which accounts for its oscillatory behaviour. If the input is too low, the activity instead converges to a stable equilibrium point, and different oscillatory bands can be obtained by varying the time constants.

The model was originally developed for visual processing and later extended to other sensory-motor and cognitive processes (Kilpatrick, 2013). Although it gives a useful account of oscillations, it has limitations that matter for the present work. Its parameters carry no direct physiological meaning; it represents only two populations, whereas a cortical region contains several distinct neuron types; it assumes a single rhythm per region; and it does not explicitly describe synaptic dynamics. The model that addresses these limitations, and that forms the basis of the unit used in this thesis, is described next.

1.3.1.b Jansen and Rit Model

The neural mass model most directly relevant to the present thesis was developed by (Jansen & Rit, 1995), building on earlier work by (Lopes da Silva, 1974). The Jansen and Rit model is also a coarse-grained or lumped parameter model which simulates a cortical column, but it improves on it in two respects: it represents the populations of a column more finely, and it gives an explicit account of synaptic dynamics, which the Wilson–Cowan model lacks. In its original form, it describes three locally interconnected populations: pyramidal neurons, excitatory interneurons using glutamate, and GABAergic inhibitory interneurons.

Each population is described by a two-stage process. First, the population receives a postsynaptic potential formed from the summed synaptic contributions of the upstream populations, and this potential is converted into a mean discharge rate through a sigmoidal function. The sigmoid transforms the postsynaptic potential, denoted y_i for postsynaptic population i , into the mean density of action potentials, or spike rate, denoted as z_i . Second, that discharge rate is converted back into a postsynaptic potential through the characteristic dynamics of the synapse, and the resulting standard potential is multiplied by the connectivity constants C_{ij} of the connection before it acts on the target population. The postsynaptic

potentials may be excitatory (EPSPs) or inhibitory (IPSPs), based on whether the synapse is glutamatergic or GABAergic.




Figure 1.2 Jansen and Rit model of three interconnected populations.

The populations are wired together by a few physiological rules. The pyramidal neurons excite both interneuron groups, the excitatory and the inhibitory. The excitatory interneurons feedback only to the pyramidal neurons, and the inhibitory interneurons act only on the pyramidal neurons. External input, carried by pyramidal neurons in other columns, can reach all three populations. The pyramidal and excitatory interneuron synapses are glutamatergic and share the same impulse response; the inhibitory interneurons are GABAergic.

Each synapse is described by a second-order dynamics. For a population using a synapse with gain G_n and time constant τ_n , the standard postsynaptic potential y_n evolves as:

$$\frac{d^2 y_n(t)}{dt^2} = \frac{G_n}{\tau_n} z_n(t) - \frac{2}{\tau_n} \frac{dy_n(t)}{dt} - \frac{y_n(t)}{\tau_n^2} \quad (1.4)$$

Where z_n is the presynaptic spike density and the subscript n denotes the population (p for pyramidal, e for excitatory interneurons, i for inhibitory interneurons). Each second-order equation is equivalent to two first-order equations,

$$\begin{cases} \frac{dy_n(t)}{dt} = x_n(t) \\ \frac{dx_n(t)}{dt} = \frac{G_n}{\tau_n} z_n(t) - \frac{2}{\tau_n} x_n(t) - \frac{y_n(t)}{\tau_n^2} \end{cases} \quad (1.5)$$

It is the form of equations that are used in simulation of present model. For each population, the membrane potential is calculated as the weighted sum of the synaptic contributions reaching

the populations. The corresponding spike rate is then obtained through the sigmoid $S(\cdot)$. Applying these steps to the three populations yields the full system of equations for the column:

Pyramidal Neurons:

$$\left\{ \begin{array}{l} v_p(t) = C_{pe}y_e(t) - C_{pi}y_i(t) \\ S(v_p(t)) = z_p(t) = \frac{2e_0}{1 + e^{r(s_0 - v_p)}} \\ \frac{dy_p(t)}{dt} = x_p(t) \\ \frac{dx_p(t)}{dt} = \frac{G_e}{\tau_e}z_p(t) - \frac{2}{\tau_e}x_p(t) - \frac{y_p(t)}{\tau_e^2} \end{array} \right. \quad (1.6)$$

Excitatory interneurons:

$$\left\{ \begin{array}{l} v_e(t) = C_{ep}y_p(t) \\ S(v_e(t)) = z_e(t) = \frac{2e_0}{1 + e^{r(s_0 - v_e)}} \\ \frac{dy_e(t)}{dt} = x_e(t) \\ \frac{dx_e(t)}{dt} = \frac{G_e}{\tau_e} \left(z_e(t) + \frac{n(t)}{C_{pe}} \right) - \frac{2}{\tau_e}x_e(t) - \frac{y_e(t)}{\tau_e^2} \end{array} \right. \quad (1.7)$$

Inhibitory interneurons:

$$\left\{ \begin{array}{l} v_i(t) = C_{ip}y_p(t) \\ S(v_i(t)) = z_i(t) = \frac{2e_0}{1 + e^{r(s_0 - v_i)}} \\ \frac{dy_i(t)}{dt} = x_i(t) \\ \frac{dx_i(t)}{dt} = \frac{G_i}{\tau_i}z_i(t) - \frac{2}{\tau_i}x_i(t) - \frac{y_i(t)}{\tau_i^2} \end{array} \right. \quad (1.8)$$

Where C_{ij} are the internal connectivity constants between populations, $n(t)$ is the external input, and $S(\cdot)$ is the sigmoid converting membrane potential to spike rate.

The model has been extended in ways that lead directly to the unit used in this thesis. (Wendling, 2002) separated the single inhibitory population into slow and fast interneurons, which broadened the range of rhythms the model could reproduce, and (David & Friston, 2003) generalized it to generate several rhythms from a single region. Building on these, (Ursino et al., 2010) developed a four-population column, with pyramidal neurons, excitatory interneurons, and slow and fast GABAergic interneurons, in which the fast inhibitory

interneurons receive a self-loop that produces robust oscillation in the gamma range. This four-population gamma column serves as the computational unit for every semantic feature in the present model; its full equations, connectivity, and parameter values are given in the methodology.

CHAPTER 2: METHODOLOGY

This chapter describes the model in two parts. The first part specifies the model this thesis starts from: the semantic-memory network of (Ursino & Pirazzini, 2024), together with the alpha-gating mechanism of (Ursino, 2025) as implemented here. The second part specifies what this thesis adds to that base: an inhibition mechanism, a spotlight mechanism, and a fatigue mechanism. The aim is to give the equations, parameters, and procedures in enough detail to reproduce the model; the motivation for the components and their place in the literature was set out in Chapter 1.

2.1 THE OLD MODEL

2.1.1 The Neural Mass Model of a Single Cortical Column

The building block of the network is a computational unit that represents a semantic feature and is modelled as a cortical column. Each column contains four populations: pyramidal neurons, excitatory interneurons (both glutamatergic), slow GABAergic inhibitory interneurons, and fast GABAergic inhibitory interneurons. This is the four-population neural mass model of (Ursino et al., 2010), introduced in section 1.3.1.b.

For each population, the mean membrane potential $v(t)$ is first computed as the weighted sum of the synaptic contributions reaching the population. The mean firing rate $z(t)$ is then obtained from $v(t)$ through a sigmoidal activation function $S(v(t))$. Finally, a normalized postsynaptic potential $y(t)$ is obtained from the synaptic dynamics; this potential is multiplied by the synaptic weight to give the actual contribution to the postsynaptic membrane potential.

Both pyramidal neurons and fast inhibitory interneurons can receive an external input, labelled u_p and u_f , respectively. Both reach the target population via excitatory synapses. In the case of u_p , a common simplification is used to reduce the number of differential equations (and therefore, the number of state variables). Specifically, we processed u_p through the excitatory synapse that goes from excitatory interneurons to pyramidal neurons (see Eq. (2.2) below), instead of processing the input separately. The other external input u_f , reaches its target population through a dedicated synapse (Eq. (2.4)).

The equations for four populations are as follows;

Pyramidal neurons:

$$\left\{ \begin{array}{l} v_p(t) = C_{pe}y_e(t) - C_{ps}y_s(t) - C_{pf}y_f(t) \\ S(v_p(t)) = z_p(t) = \frac{2e_0}{1 + e^{r(s_0 - v_p)}} \\ \frac{dy_p(t)}{dt} = x_p(t) \\ \frac{dx_p(t)}{dt} = \frac{G_e}{\tau_e}z_p(t) - \frac{2}{\tau_e}x_p(t) - \frac{y_p(t)}{\tau_e^2} \end{array} \right. \quad (2.1)$$

Excitatory interneurons:

$$\left\{ \begin{array}{l} v_e(t) = C_{ep}y_p(t) \\ S(v_e(t)) = z_e(t) = \frac{2e_0}{1 + e^{r(s_0 - v_e)}} \\ \frac{dy_e(t)}{dt} = x_e(t) \\ \frac{dx_e(t)}{dt} = \frac{G_e}{\tau_e} \left(z_e(t) + \frac{u_p}{C_{pe}} \right) - \frac{2}{\tau_e}x_e(t) - \frac{y_e(t)}{\tau_e^2} \end{array} \right. \quad (2.2)$$

Slow inhibitory interneurons:

$$\left\{ \begin{array}{l} v_s(t) = C_{sp}y_p(t) \\ S(v_s(t)) = z_s(t) = \frac{2e_0}{1 + e^{r(s_0 - v_s)}} \\ \frac{dy_s(t)}{dt} = x_s(t) \\ \frac{dx_s(t)}{dt} = \frac{G_s}{\tau_s}z_s(t) - \frac{2}{\tau_s}x_s(t) - \frac{y_s(t)}{\tau_s^2} \end{array} \right. \quad (2.3)$$

Fast inhibitory interneurons:

$$\left\{ \begin{array}{l} v_f(t) = C_{fp}y_p(t) - C_{fs}y_s(t) - C_{ff}y_f(t) + y_l(t) \\ S(v_f(t)) = z_f(t) = \frac{2e_0}{1 + e^{r(s_0 - v_f(t))}} \\ \frac{dy_f(t)}{dt} = x_f(t) \\ \frac{dx_f(t)}{dt} = \frac{G_f}{\tau_f}z_f(t) - \frac{2}{\tau_f}x_f(t) - \frac{y_f(t)}{\tau_f^2} \\ \frac{dy_l(t)}{dt} = x_l(t) \\ \frac{dx_l(t)}{dt} = \frac{G_e}{\tau_e}(u_f(t)) - \frac{2}{\tau_e}x_l(t) - \frac{y_l(t)}{\tau_e^2} \end{array} \right. \quad (2.4)$$

Where the subscript l is used to represent the quantities in the additional synapse, introduced to describe the effect of the input u_f via glutamatergic dynamics. All related parameters' values are listed in the tables below. Note that the variance of the noise is divided by the integration step, dt . In this way, we obtain a white noise with power density equal to $\sigma^2 dt = 5$.

Synapses				Function $S(v(t))$	
G_e (mV)	5.17	τ_e (ms)	8	e_0 (Hz)	2.5
G_s (mV)	4.45	τ_s (ms)	33.4	r (mV ⁻¹)	0.56
G_f (mV)	57.1	τ_f (ms)	2.5	s_0 (mV)	12

Table 2.1 Parametric values for Synapses and Functions

Inputs		Intra-column connections			
m_p (Hz)	500 or 0	C_{ep}	54	C_{fp}	108
m_f (Hz)	500 or 0	C_{pe}	54	C_{fs}	27
σ_p^2 (s ⁻²)	5/dt	C_{sp}	54	C_{pf}	300
σ_f^2 (s ⁻²)	5/dt	C_{ps}	67.5	C_{ff}	10

Table 2.2 Parametric values for inputs and connectivity constants

Each unit receives external input through two pathways, one to the pyramidal/excitatory side and one to the fast inhibitory side. The input on each pathway combines Gaussian white noise with the other signals reaching that population. The two noise processes, $n_p(t)$ and $n_f(t)$, are independent, with zero mean and standard deviation $\sigma = \sqrt{5/dt}$, where the $1/dt$ factor keeps the power spectral density constant across integration steps and yields a white-noise process with power density $\sigma^2 \cdot dt = 5$. The noise n_p is injected into the excitatory-interneuron pathway, exciting the pyramidal population indirectly, while n_f is injected into the fast inhibitory population through the additional state variable y_l .

The total input to each pathway is then

$$u_p(t) = n_p(t) + m(t) + \sum W_{ij}^p z_{p,j}(t - D_{ij}) \quad (2.5)$$

$$u_f(t) = n_f(t) + \sum W_{ij}^f z_{p,j}(t - D_{ij}) \quad (2.6)$$

Where $m(t)$ is the external stimulus, non-zero only for a presented feature, and the summation terms are the recurrent contributions from the trained connections, the excitatory weights W_{ij}^p and the inhibitory weights W_{ij}^f acting on the presynaptic pyramidal spike densities $z_{p,j}$ with

conduction delay, D_{ij} . For u_p a common simplification is used to reduce the number of state variables: rather than processing it through a separate synapse, it is passed through the excitatory synapse from the excitatory interneurons to the pyramidal neurons (Eq. 2.2). The input u_f reaches its target through a dedicated synapse (Eq. 2.4). The noise introduces trial-to-trial variability in the timing of the gamma oscillations, so that units coding for different features begin with independent phases; this variability is the basis of the seed-dependence examined in the sensitivity analysis of Chapter 3.

2.1.2 The Animal Taxonomy

The semantic content of the model is a set of 102 features describing 14 animals, drawn from the taxonomy of (Ursino & Pirazzini, 2024). The animals fall into two categories: 8 mammals (dog, cat, sheep, cow, giraffe, zebra, bear, and lion) and 6 birds (rooster, hen, owl, pigeon, goose, and parrot). Each animal is described by features that capture its salient and marginal properties, as well as the properties it shares with, or that distinguish it from, other members of its categories. The 102 features are indexed from 1 to 102, and each is represented by one neural mass unit in the semantic store. The full feature list is given in Table 2.3, and the taxonomy is shown in Figure 2.1.

The taxonomy is organized hierarchically. At the most general level, three features are shared by all 14 animals: "it eats" (100), "it drinks" (101), and "it sleeps" (102), and these define the superordinate category *animal*. One level down, the mammals share "it has four legs" (58), "it has fur" (59), and "it is viviparous" (60), while the birds share "it has two legs" (97), "it has feathers" (98), and "it has wings" (99). Within each category, further subcategories are defined by additional shared features: the herbivorous mammals (sheep, cow, giraffe, zebra) share "it eats grass" (41) and "it is a ruminant" (42); the dangerous animals (bear, lion) share "it is dangerous" (54) and "it has big claws" (55); and the flying birds (owl, pigeon, goose, parrot) share "it flies" (95) and "it makes the nest" (96). These subcategories can intersect, as shown in Figure 2.1.

Each animal is further described by its own distinctive features, which may be salient or marginal. The lion, for instance, has five salient features ("it roars", 48; "it has a bushy mane", 49; "it is brave", 50; "it is a predator", 51; "it has sharp teeth", 52) and two marginal ones ("it lives in herds", 40, shared with the zebra; "it defends the territory", 53). The distinction between salient and marginal features is not assigned by hand but arises from how often a feature is presented during training: features presented rarely become marginal and are not spontaneously evoked when the concept is recalled.

A few features are shared between pairs of animals without belonging to any labelled subcategory. "It has spots" (28) is shared by the cow and the giraffe, and "it pecks" (65) by the

rooster and the hen. These pairwise shared features are a useful test case, since they lack a consistent category label and sit outside the explicit hierarchy.

Category / Animal	No.	Feature	Type
Dog	1	It barks	Salient
	2	It wags its tail	
	3	It is loyal	
	4	It keeps watch	
	5	It is affectionate	
	6	It sleeps in the kennel	Marginal
	7	It is of various sizes	
Cat	8	It meows	Salient
	9	It purrs	
	10	It is independent	
	11	It chases mice	
	12	It has whiskers	Marginal
	13	It scratches	
Domestic	14	It is domestic	Salient
	15	It is a pet	
Sheep	16	It bleats	Salient
	17	It is sheared	
	18	It produces wool	
	19	It lives in the flock	
	20	It is soft	Marginal
	21	It has curly hair	
Cow	22	It moos	Salient
	23	It has udders	
	24	It has horns	
	25	It lives in the stable	
	26	It has a bell	Marginal
	27	It is sacred in India	
	28	It has spots*	
Farm	29	It produces milk	Salient
	30	It lives in the farm	
Giraffe	31	It has a long neck	Salient
	32	It has long legs	
	33	It is very high	
	28	It has spots*	Marginal

	34	It has thin legs	
	35	It has a prehensile tongue	
Zebra	36	It has stripes	Salient
	37	It runs	
	38	It has hooves	
	39	It camouflages itself	Marginal
	40	It lives in herds**	
Herbivorous	41	It eats grass	Salient
	42	It is a ruminant	
Bear	43	It is omnivorous	Salient
	44	It goes into hibernation	
	45	It eats honey	
	46	It is big	Marginal
	47	It lives in the mountain	
Lion	48	It roars	Salient
	49	It has a bushy mane	
	50	It is brave	
	51	It is a predator	
	52	It has sharp teeth	
	40	It lives in herds**	Marginal
	53	It defends the territory	
Dangerous	54	It is dangerous	Salient
	55	It has big claws	
Wild	56	It is wild	Salient
	57	It lives in nature	
Mammals	58	It has four legs	Salient
	59	It has fur	
	60	It is viviparous	
Rooster	61	It crows at dawn	Salient
	62	It lives in the coop	
	63	It has a big crest	
	64	It has a spur	
	65	It pecks***	Marginal
	66	It is aggressive	
Hen	67	It makes chicks	Salient

	68	It hatches eggs	
	69	It has a small crest	
	65	It pecks***	Marginal
	70	It is stupid	
Owl	71	It lives at night	Salient
	72	It hunts rodents	
	73	It is a raptor	
	74	It has a rostrum beak	
	75	It has keen sight	
	76	It lives in the woods	Marginal
	77	It is silent	
Pigeon	78	It coos	Salient
	79	It eats the crumbs	
	80	It is in the squares	
	81	It litters	
	82	It is used to send messages	Marginal
	83	It carries diseases	
Goose	84	It squawks	Salient
	85	It has webbed feet	
	86	It has an orange beak	
	87	It lives in the pond	Marginal
	88	It migrates	
Parrot	89	It has a big beak	Salient
	90	It repeats words	
	91	It speaks	
	92	It has a colored tail	
	93	It is exotic	Marginal
	94	It lives as a couple	
Volatile	95	It flies	Salient
	96	It makes the nest	
Birds	97	It has two legs	Salient
	98	It has feathers	
	99	It has wings	
Animals	100	It eats	Salient
	101	It drinks	
	102	It sleeps	

Table 2.3 Semantic features of 14-animals taxonomy: Adapted from (Ursino & Pirazzini, 2024). Features marked * are shared by the cow and the giraffe, ** by the zebra and the lion, and *** by the rooster and the hen.




Figure 2.1 Taxonomy of animals (Ursino, 2024)

2.1.3 Network Training

The network is trained with the method of (Ursino & Pirazzini, 2024), without modification. The connections between units are learned through an asymmetric, presynaptically gated variant of the Hebb rule.

Two types of connections are trained, both starting from zero. The excitatory connection W_{ij}^{ex} runs from the pyramidal population of presynaptic unit j to the pyramidal population of postsynaptic unit i . The inhibitory connection W_{ij}^{in} is disynaptic, running from the pyramidal population of unit j to the fast inhibitory population of unit i , which, in turn, inhibits the pyramidal population of unit i .

The activities used by the learning rule are obtained as follows. The spike density of each population is computed from its mean membrane potential through a sigmoidal activation function:

$$z(t) = S(v(t)) = \frac{2e0}{1 + e^{-r(v-s0)}} \quad (2.7)$$

Here $v(t)$ is the mean membrane potential, e the half-maximum firing rate, r the slope of the sigmoid, and $s0$ the potential at which the population reaches half its maximum rate. The function takes the form above and saturates at $2e0$, with every population's firing rate falling between 0 and $2e0$. This same maximum rate sets the scale for learning. Before the Hebb rule is applied, each spike density is averaged over the previous 30ms and divided by $2e0$, which maps it onto the 0 -1 range, and the result is compared with the thresholds.

The rule is made asymmetric by a presynaptic gate: an update is applied only when the presynaptic activity is above its threshold. Below the thresholds, the connection does not change, whatever the postsynaptic unit is doing. When the gate is open, the postsynaptic side sets the sign: activity above threshold potentiates the connection, activity below it depresses the connection. The change in the excitatory connection from unit j to unit i is:

$$\Delta W_{ij}^{ex} = \gamma^{ex} (\bar{z}_{pi}(t) - \theta_{post}^{ex}) (\bar{z}_{pj}(t) - \theta_{pre}^{ex})^+ (W_{max}^{ex} - W_{ij}^{ex}(t)) \quad (2.8)$$

Where $\bar{z}_{p,i}$ and $\bar{z}_{p,j}$ are the averaged pyramidal spike densities of the postsynaptic and presynaptic units, θ_{post}^{ex} and θ_{pre}^{ex} are the postsynaptic and presynaptic thresholds respectively, and γ^{ex} is the learning rate. The notation $()^+$ is the positive part, $(u)^+ = u$ if $u > 0$ and 0 otherwise, so the presynaptic factor acts as the on/off gate. The final factor $(W_{max}^{ex} - W_{ij}^{ex})$ slows learning as the connection nears its saturation level. The inhibitory connection follows the same form, with its postsynaptic term taken from the averaged fast-interneuron activity $\bar{z}_{f,i}$ and its own learning rate and thresholds:

$$\Delta W_{ij}^{in} = \gamma^{in} (\bar{z}_{f,i}(t) - \theta_{post}^{in}) (\bar{z}_{p,j}(t) - \theta_{pre}^{in})^+ (W_{max}^{in} - W_{ij}^{in}(t)) \quad (2.9)$$

This gating is what builds the hierarchy. The intended connectivity reflects two distinctions: between shared and distinctive features, and between salient and marginal features. A distinctive feature should send connections up to the shared features of its object but should not receive them, and a marginal feature should send connections to the salient features it helps recall but should not be driven by them. The gate produces this directionality. When a distinctive feature and a shared feature are active together, both connections between them strengthen. When the shared feature later appears with a different animal while the distinctive feature is silent, only the shared-to-distinctive connection weakens, because the distinctive unit is not presynaptically active to protect its outgoing link. Over many presentations, connections from shared features down to distinctive ones fade, while connections from distinctive features up to shared ones survive. The same frequency logic separates salient from marginal features.

Training is probabilistic. In each epoch, the fourteen animals are presented once in random order, and within a presentation each feature is switched on with a fixed probability: 80% for salient and shared features, 30% for marginal features. An active feature receives an input of 500, enough to drive its unit into the gamma band without saturating it. Each presentation runs for 0.35 s, and the rule is applied during the last 0.25 s, after the initial transient. After each weight update, the self-connection of every unit is constrained to remain zero ($W_{jj}=0$), so a feature cannot reinforce a synapse onto itself; It is worth noticing that self-connections in this model are always zero. Further normalization has been adopted for the sum of synapses entering a given unit, i.e., we assumed this sum cannot overcome a given upper saturation. This constraint reproduces a physiological limitation in overall neurotransmitter availability and is helpful to avoid excessive excitation. This matters most for the heavily shared category units,

which would otherwise saturate and lose their gamma oscillation. The network is trained for 600 epochs, by which point the weights are close to stable.

Parameter	Value
dt (ms)	1
D (ms)	25
θ_{post_ex}	0.25
θ_{post_in}	0.05
θ_{pre_ex}	0.08
θ_{pre_in}	0.16
γ_{ex}	0.00009
γ_{in}	0.000024
W_{max}^{ex}	460
W_{max}^{in}	70
W_{maxsum}^{ex}	2300
W_{maxsum}^{in}	280

Table 2.4 Values of training parameters

After training, a single feature presented as a cue drives the network to recover the full set of salient features for that animal and its categories, synchronized in the gamma band, while marginal and unrelated features stay silent. This trained network is the semantic store on which the new model in Section 2.2 is built.

2.2 NEW MODEL

The new model places the trained semantic store of Section 2.1.2 inside a three-region architecture and adds three mechanisms that act on it: an inhibition mechanism, a moving spotlight, and synaptic fatigue. This section first specifies the three regions and their connections, then describes each mechanism in detail. The motivation for these additions was given in Chapter 1; what follows is their implementation.

2.2.1 Multi-region architecture

The model has three cortical regions (ROIs), each is constituted by a network of neural mass units joined by long-range pyramidal connections: ROI 1 (semantic store), ROI 3 (alpha oscillator), and ROI 4 (recall area). The labels are consistent with Ursino (2025), and notably, there is no ROI 2 in this network.

ROI 1 — Semantic store: ROI 1 holds the 102 feature units of the animal taxonomy (Section 2.1.3), each built from the same four-population gamma column, one unit per feature, with the connectivity, time constants, and gains given in Table 1. Every unit stands for a single semantic

feature and receives three inputs: the external stimulus, background noise, and the alpha rhythm projected from ROI 3, which gates the unit open or closed to attention. Once a unit is released by the alpha gate, its activity passes forward through its pyramidal neurons to ROI 4. ROI 1 is where the individual features are represented and where the attentional spotlight acts, selecting which feature is released for downstream processing at any given moment.

ROI 3 — Alpha oscillator: ROI 3 is a single neural mass unit tuned to oscillate at about 10 Hz. Two connectivity changes distinguish it from the gamma units: the slow inhibitory connectivity is much larger $C_{ps} = 450$ (against 67.5), and the fast inhibitory connectivity is reduced $C_{fp} = 35$ (against 108), which together weaken the gamma-generating mechanism and pull the frequency into the alpha band. Its time constants are: $\tau_e = 15.2$, $\tau_s = 23.8$, $\tau_f = 3.3$ ms. The constant drive is $m_3 = 1000$. The pyramidal spike density $z_{p3}(t)$ projects to the fast GABAergic populations of both ROI 1 and ROI 4. The projection to ROI 4 uses $W_{43}^f = 100$ with a delay; the projection to ROI 1 is gated by attention, as described above.

ROI 4 — Recall area: ROI 4 mirrors ROI 1 in structure, 102 gamma-band units, one per feature, using the same four-population column and the same parameters, but it receives no external stimulus of its own. Its input comes from two sources. The first is a delayed excitatory projection from ROI 1, carried by the weight $W_{41}^p = 200$, together with an inhibitory projection from the alpha oscillator in ROI 3, carried by $W_{43}^p = 100$; both arrive with a negligible delay ($D_4 = 0.0001$ s). The second is ROI 4's own recurrent excitatory connectivity: the trained matrix, scaled by the depletion variable of each unit, feeds back onto ROI 4's units with a 30 ms delay. This recurrence is what reconstructs a full concept from the single feature relayed by ROI 1, filling in the salient and shared features associated with the attended animal. Because the recurrent weights are continuously scaled by depletion, a concept's self-sustaining activity weakens the longer it is held active, allowing ROI 4 to release one retrieval and complete the next as attention moves.




Figure 2.2 Three-region architecture of the model

2.2.2 Inhibition Mechanism

Inhibitory control of the semantic store is exerted top-down by the alpha oscillator. The alpha rhythm generated in ROI 3 reaches the fast GABAergic population of each unit in ROI 1 through the connection W_{13}^f , and this projection determines whether a feature can pass to the recall area. By default, $W_{13}^f = 300$ for every unit, so the alpha rhythm suppresses the entire store: in this condition, each feature is held under rhythmic inhibition and cannot drive ROI 4.

The suppression works by phase opposition. The delay $D_1 = 50\text{ms}$ places the alpha rhythm arriving at ROI 1 in approximate phase opposition (about 180°) to the rhythm arriving at ROI 4 through W_{43}^f . A suppressed feature is therefore maximally active precisely when ROI 4 is maximally inhibited, so its signal cannot be read out downstream. With $D_1 = 50\text{ms}$ and an alpha period $T_\alpha \approx 100\text{ms}$, this 180° condition is met, satisfying the phase requirement for gating identified by (Ursino, 2025). Because the projection targets the fast inhibitory interneurons, the gate acts through inhibition; routing it to the pyramidal population instead would not suppress the feature.

2.2.3 Spotlight Mechanism

Attention is implemented as a selective release of this inhibition. To attend to a feature, its connection W_{13}^f is set to zero, which removes the alpha suppression from that unit alone while the remaining 101 units stay gated. The released feature can then reach the recall area and to drive retrieval in ROI 4, so a single cued feature is sufficient to recall its associated concept while all other concepts remain suppressed.

Moving the released window from one feature to another redirects attention across the store, producing a spotlight that scans the network in time. In the present simulations, the gate is opened on one cue feature at a time over a single run: suppression is lifted from the first animal's cue during 0–0.4 s, to the second during 0.8–1.2 s, and to the third during 1.6–2.0 s, with the rest of the store held inhibited throughout, so that attention visits the three concepts in sequence. The triplet is not fixed; any three features can be chosen by selecting which connections are released in which window. Because suppression is governed by a single incoming connection per unit, shifting attention requires only deciding which connection is released during each temporal window.

2.2.4 Fatigue Mechanism

A use-dependent synaptic fatigue is introduced in the recall area so that an active concept cannot dominate indefinitely, and the network can disengage from it without an external

command. The fatigue is modelled as depletion of the available transmitter pool at the recurrent synapses within ROI 4. Each of the 102 units carries a scalar depletion variable $D_i(t) \in [0,1]$, governed by

$$\frac{dD_i(t)}{dt} = \alpha(1 - D_i(t)) - \beta z_{p4,i}(t)D_i(t) \quad (2.10)$$

Where $z_{p4,i}$ is the unit's pyramidal spike density, $\alpha = \frac{1}{\tau_d}$ is the recovery rate toward full availability, with $\tau_d = 1$ s, and $\beta = 2\alpha$ is the depletion rate. The first term restores D_i toward 1 when the unit is silent; the second depletes it at a rate proportional to the current firing.

The depletion scales the recurrent excitatory connectivity of the recall area row-wise: the trained excitatory matrix is multiplied by D_i on the synapses entering unit i . This produces a new, depletion-scaled matrix, written $W_{depleted}^p$, and it is this matrix, not the original trained one, that drives ROI 4's recurrent excitation. Units that have fired strongly have a depleted pool ($D_i \rightarrow 0$) and receive reduced recurrent drive, while silent units retain a full pool ($D_i \rightarrow 1$) and remain readily activatable. Since the recurrent inhibition is off, fatigue acts on the excitatory recurrence alone.

This asymmetry is what drives the sequential retrieval. Once a concept's features have stayed active for roughly 400 ms, their depletion variables have dropped far enough to weaken the recurrent drive that sustains them; the representation fades, and the recall area is free for the next concept as the spotlight shifts. Each depletion variable starts at $D_i(0)=1$ and is updated every step by forward Euler. The coupling at each step is one-directional. Depletion scales the effective recurrent weights, the altered weights change the membrane potentials and spike densities, and these feed back into the depletion at the next step. This feedback runs on a timescale of hundreds of milliseconds, far slower than the gamma oscillations (~ 30 ms period) or the alpha rhythm (~ 100 ms period), so it controls the turnover between concepts without disturbing the dynamics within a single retrieval.

CHAPTER 3: SIMULATIONS AND RESULTS

The model is evaluated in two stages. First, a sensitivity analysis tests how the recall behaviour depends on the random seed and identifies which parameters need tuning to keep the results clean. Second, using the parameter set fixed by that analysis, the spotlight mechanism is applied to a range of animal triplets to check that the network recalls each attended concept correctly as attention moves across it.

Throughout this chapter, recall is assessed against the expected oscillation pattern of the trained network. In each spike-density plot, the line colours denote feature types as follows: a solid black line for its salient features and a dotted black line for its marginal features. The category lines mark its inherited subcategory features, shown as a dash-dot blue line for "flying bird" and a dashed cyan line for "bird". The dashed-dot green line marks the "animal" features shared by all animals, and the solid yellow line marks the "other features" outside the attended animal's categories, which are expected to stay silent. A retrieval is judged correct when the attended animal's salient features and all its inherited shared-feature lines, including the green line, are evoked in synchrony, and a failure when one or more of these shared lines is absent or when the yellow line is evoked.

Each animal inherits a different set of subcategories, so its category lines and colors are not the same from one figure to the next; the legend attached to each graph gives the exact mapping for that case.

3.1 SENSITIVITY ANALYSIS AND PARAMETER VALIDATION

Because each unit receives independent Gaussian noise, the network's behaviour depends on the random seed. To test the robustness, eight random seeds were used, rng(10) through rng(80), with the model evaluated on one representative animal triplet: (Cat-Hen-Sheep). The parameters examined were the recurrent excitation weight W_p (multiplied through a scaling factor), the depletion closing-gate parameter β , the fatigue recovery time constant τ_d , and the standard deviation of the Gaussian white noise (SD). This analysis, and all simulations that follow, use a fixed spotlight schedule over a single 2.4 s run: the gate is opened on the first animal's cue feature during 0.0–0.4 s, on the second during 0.8–1.2 s, and on the third during 1.6–2.0 s, with the rest of the store held under alpha suppression throughout, including the 0.4 s gaps between windows.

The results are summarized in Table 3.1. Different values of β , τ_d , and SD were tested during preliminary tuning, and $\beta=2\alpha$, $\tau_d=1.0s$, and SD =5 gave the best and most consistent retrieval across all eight seeds; these values were then held fixed for the rest of the analysis. The parameter that carries the behaviour is the recurrent excitation scaling W_p . A value of $W_p = W_p \times 2.1$ produced correct retrieval for most cases, with two exceptions. At rng(50), the gamma

rhythm of the first attended animal, "Cat" (spotlight 0.0–0.4 s), takes almost 0.8 s to return to baseline. As a result, its features carry over and evoke during the second window, dedicated to "Hen" (0.8–1.2 s). This is corrected either by reducing W_p to 2.0 or by setting $SD = 3$; the W_p change is adopted, using 2.0 for this case. Secondly, at $\text{rng}(30)$, the model does not retrieve the shared features within the second window ("Hen", 0.8–1.2 s), although the shared features do appear shortly after 1.2 s, which indicates that in this case the network needs more time to recover the shared features from the other animals. Increasing W_p slightly to 2.3 retrieves the shared features within the window dedicated to Hen. These adjustments were made on a per-seed basis, and in every case a stable parameter configuration was found that restored correct behaviour.

parameters	$\text{rng}(10)$	$\text{rng}(20)$	$\text{rng}(30)$	$\text{rng}(40)$	$\text{rng}(50)$	$\text{rng}(60)$	$\text{rng}(70)$	$\text{rng}(80)$
W_p	2.1	2.1	2.3	2.1	2.0	2.1	2.1	2.1
β	2α	2α	2α	2α	2α	2α	2α	2α
τ_d	1	1	1	1	1	1	1	1
SD	5	5	5	5	5	5	5	5

Table 3.1 Sensitivity of the model parameters to the random seed.

The standard parameter set used for all subsequent simulations is therefore: $W_p = W_p \times 2.1$, $\beta = 2\alpha$, $\tau_d = 1.0$ s, and $SD = 5$. This combination produced consistent retrieval across the wide range of tested conditions and is the basis for all results reported in Table 3.2.

3.2 SELECTIVE ATTENTION – SPOTLIGHT MECHANISM

The selective-attention mechanism is applied to triplets of animals drawn from different categories, with each animal attended during the spotlight schedule introduced in Section 3.1. The triplet is not fixed: any three concepts can be selected by choosing which three cue features the gate releases, in sequence. For each triplet, recall is checked under the same eight seeds ($\text{rng}(10)$ through $\text{rng}(80)$). A result is marked correct ("ok") when the attended animal evokes its full expected pattern with all oscillations present, and a failure when the features shared by all animals are not evoked, that is, when the green oscillations are absent (Table 3.2).

The selective-attention mechanism is applied to triplets of animals drawn from different categories, with each animal attended during the spotlight schedule introduced in Section 3.1. The triplet is not fixed: any three concepts can be selected by choosing which three cue features the gate releases, in sequence. For each triplet, recall is checked under the same eight seeds ($\text{rng}(10)$ through $\text{rng}(80)$). A result is marked correct ("ok") when the attended animal evokes its full expected pattern with all oscillations present, and a failure when the features shared by all animals are not evoked, that is, when the green oscillations are absent (Table 3.2). In every case tested, the attended animal's own distinctive features come out correctly and are never what causes a failure. What changes between "ok" and "fail" is whether the shared, all-animal

features also appear and whether the features outside the animal's categories stay suppressed, as the yellow line should.

Animals	rng(10)	rng(20)	rng(30)	rng(40)	rng(50)	rng(60)	rng(70)	rng(80)
Cat	Ok	Ok	Ok	Ok	Fail	Ok	Ok	Ok
Hen	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok
Sheep	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok
Dog	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok
Roster	Ok	Ok	Ok	Ok	Ok	Fail*	Ok	Ok
Owl	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok
Giraffe	Fail*	Fail*	Fail*	Fail*	Fail*	Fail*	Fail*	Fail*
Cow	Ok	Ok	Ok	Ok	Ok	Ok	Fail	Ok
Pigeon	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok
Goose	Ok	Ok	Fail	Ok	Ok	Ok	Ok	Ok
Bear	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok
Parrot	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok
Giraffe	Fail*	Fail*	Fail*	Fail*	Fail*	Fail*	Fail*	Fail*
Zebra	Fail*	Fail*	Fail*	Fail*	Fail*	Fail*	Fail*	Fail*
Lion	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok
Owl	Ok	Ok	Fail	Ok	Ok	Ok	Ok	Ok
Pigeon	Ok	Ok	Ok	Ok	Ok	Ok	Fail	Ok
Parrot	Ok	Ok	Ok	Ok	Ok	Ok	Fail*	Ok
Hen	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok
Goose	Ok	Ok	Ok	Ok	Ok	Ok	Ok	Ok
Rooster	Fail*	Ok	Fail*	Fail*	Fail*	Fail*	Fail*	Fail*


Table 3.2: Each cell reports the outcome of the spotlight recall for that triplet under that seed, judged visually against the expected oscillation pattern. ‘ok’ = all features recalled correctly, with every oscillation pattern as expected. ‘fail’ = unexpected activation of the yellow lines (“other features” that should remain silent). ‘fail*’ = the shared features were not evoked (the green, all-animal lines are absent).

Of the 56 triplet–seed combinations tested, the large majority recall correctly under the standard parameters, a smaller number are recovered by a per-case increase in W_p . Only the giraffe–zebra–lion combination at certain seeds remains unrecoverable without introducing saturation. This confirms that the spotlight generalizes: the mechanism works for arbitrary groups of three animals, not only for one fixed set. The failures that occur are isolated rather

than systematic; they appear at particular seed–triplet combinations, usually as one attended animal in the triplet losing its shared lines while the other two recall correctly. The remainder of this section presents representative visualizations from the model.

3.2.1 A fully correct case

The Dog–Rooster–Owl triplet at rng(10) shows all three animals behaving correctly: each attended animal evokes its salient and shared features within its window, and the others stay suppressed. Furthermore, note that, at the end of the correct period, the features disappear thanks to the fatigue mechanisms.






Figure 3.1. A fully correct case; (A) Simulation of the network in response to a salient distinctive feature of the Dog (“it barks”). The plot represents the response of all units in the net in the time interval 0.0 – 0.8 s. Only the Dog’s salient features are evoked, distinctive and shared, and oscillate in the gamma range. In particular, the continuous black lines represent the salient distinctive features of the Dog, the dashed-dotted lines represent the features of domestic animals, the dotted magenta lines shows the features of mammals and the green dash-dotted lines the features shared by all animals and and the solid yellow line marks the “other features” outside the attended animal’s categories (remained at zero). All these features’ activities are synchronized. The residual features in the net remain all at zero and are not shown in the legend for clearness. (B) Zoomed version of last gamma cycle of Dog’s features. (C and D): correctly evoked features of Rooster and Owl with respective colored legends.

The same triplet at rng(60) shows a failure in the Rooster window, where the shared features are not evoked under the standard parameters. Raising W_p slightly to 2.2 restores them.




Figure 3.2. (A) Absence of shared features in Rooster’s window at $\text{rng}(60)$; (B) After slightly increasing the parameter, W_p to 2.2, shared features evoked.

Moreover, in some cases a yellow line is briefly active at the start of a window, which means the previous animal has not yet returned to baseline and is weakly active, in the current window. These cases are marked as failures in the table, and each is recoverable by the same small adjustment to W_p . The Cat–Hen case at $\text{rng}(50)$ is one example.





Figure 3.3. The features of cat (0.0-0.4s) has not returned to baseline within the expected window and emerged as active yellow line ('other features') in second window for Hen (0.8-1.2s)

3.2.2 A difficult concept: the giraffe (Giraffe–Cow–Pigeon)

The giraffe is a difficult case for the model. Under the standard parameters, none of its shared features are evoked at any seed, and only its salient features appear. Recovering the shared rhythms requires raising W_p to the range 2.7–3.2. Even then, the model may need more time than the allotted spotlight window to produce all the expected oscillations. The figure below, obtained at $\text{rng}(80)$ with $W_p=2.7$, shows that the giraffe's features still fail to appear within its spotlight window (0.0–0.4 s) despite the increased W_p .






Figure 3.4. Recall of the giraffe's features (Giraffe–Cow–Pigeon triplet, $\text{rng}(80)$) at two values of the recurrent excitation. **(A)** Giraffe window at standard $W_p=2.1$. The giraffe's salient features (black) are evoked, but its shared features, herbivorous, wild, mammal, and the "animal" line shared by all animals, fail to appear within the spotlight window (0.0–0.4 s). **(B)**: At $W_p=2.7$, giraffe's shared features begin to appear, but only after the spotlight window has closed (from around 0.4 s onward), confirming that the giraffe needs more time than the 0.4 s window allows even at the higher excitation level. **Line key**: black = salient giraffe features; coloured lines = inherited category features (herbivorous, wild, mammal); green = shared "animal"; the flat yellow line at zero is "other features," which correctly stays silent.


3.2.3 Competing demands: the Giraffe–Zebra–Lion triplet

The Giraffe–Zebra–Lion triplet is the hardest case in the study, because all three animals place unusually heavy demands on the recurrent network, and they do so in conflicting ways. The giraffe and zebra both carry large sets of shared-feature memberships: each belongs to the mammals and to the herbivorous subcategory (and, for the zebra, the wild subcategory), and the giraffe additionally shares "it has spots" with the cow. Completing either concept, therefore, requires the recurrent network to evoke a large and entangled set of shared features, which is why neither is recovered by the standard excitation, and both need an elevated W_p . This echoes Ursino and Pirazzini (2024), where "it has spots" was the hardest case for the trained network and required a 30% increase in excitatory strength.


The lion, however, pulls in the opposite direction. It is a comparatively well-separated concept that completes easily, so the high W_p needed to recover the giraffe and zebra drives the lion into saturation. The triplet thus combines two concepts that demand strong excitation with one that cannot tolerate it, and no single value of W_p satisfies all three at once: a value low enough to keep the lion stable leaves the giraffe and zebra incomplete, while a value high enough to complete them saturates the lion. The difficulty is therefore not a fault in the spotlight

mechanism but a property of this particular grouping, in which the recall demands of the attended concepts are mutually incompatible under a shared excitation setting.


To test whether the recurrent inhibition could resolve this trade-off, the triplet was run at the elevated $W_p=3.2$ (held fixed across seeds for this group) while the inhibitory scaling was varied across $W_f=0, 0.5, 1$, and 2 . Raising W_f progressively reduced the lion's saturation, but it did so at the cost of the zebra: as the inhibition increased, the zebra's shared features were suppressed in its own window. No value of W_f both removed the saturation and preserved the zebra's features, so the trade-off seen with W_p persists under inhibitory control rather than being resolved by it. The following case is presented at $\text{rng}(10)$ at $W_p = W_p^*3.2$:



(C) Wf=0.5



(D) Wf=1



(E) Wf = 2




Figure 3.5. Recall in the Giraffe-Zebra-Lion triplet at $W_p=3.2$, shown for the recurrent inhibition scaling $W_f=0, 0.5, 1$, and 2 . (A) Shows the plot of Lion's window, where saturation has occurred due to increased value of W_p . (B) zoomed version of same plot. (C, D, & E) represents Lion's window with $W_f=0.5, 1, 2$ respectively. (F) shows suppressed features of Zebra when W_f is increased for reducing saturation in Lions window.

The behaviour of this triplet across seeds reflects the same tension. At rng(30) the zebra's shared features cannot be recovered at all. At rng(40), $W_p=3$ already produces saturation in the lion. At rng(50), W_p in the range 3.2–3.5 fails to evoke the shared features and also saturates. At rng(60), $W_p=3.2$ leaves the zebra without shared features while $W_p=3.5$ saturates, and rng(70) behaves likewise.

3.2.4 Recall from a marginal feature

A further test examined recall when the spotlight is opened on a marginal feature rather than a salient one. In the Owl-Pigeon-Parrot triplet, the pigeon and parrot were cued through a salient

feature as before, while the owl was cued through its marginal feature "it lives in the woods." At rng(20), the owl recovered its full pattern within the spotlight window: the salient features, the flying-bird and bird category features, and the shared all-animal features were all evoked and synchronized by the end of the window (Figure 3.6).

Recall, however, develops with characteristic timing. The cued marginal feature fires first and in isolation, early in the window, before the rest of the pattern has formed. The salient features follow, and the shared, category-level features (flying bird, bird, and the green animal features) appear and lock into synchrony only later, from about 0.25 s onward. This staged build-up matches the behaviour of the trained network reported by (Ursino & Pirazzini, 2024), that a marginal cue still recovers the full concept, but more slowly than a salient cue, because the asymmetric Hebb rule leaves marginal features with weak incoming connections while preserving their outgoing connections to the salient features. The marginal feature therefore drives the concept forward, but the synchronization of the salient and shared features takes longer to establish than it would from a salient cue. The pigeon and parrot, cued through salient features, recalled normally in their windows.




Figure 3.6 Recall from the marginal feature ('it lives in the wood') of Owl. Here, the cued marginal feature fired initially, followed by the salient features that lock into synchrony after 0.25s.

3.2.5 Recovering the rooster window: Hen–Goose–Rooster

For the Hen–Goose–Rooster triplet, the rooster window needs a raised W_p to recover the shared-feature peaks, with the required value differing by seed. At rng(10), rng(30) and rng(40), a value of $W_p=2.3$ recovers the rooster window. At rng(50) and rng(80) the value must be raised to 2.5. At rng(60) and rng(70), the value must be raised to 2.7. In each case, the adjustment is local to the rooster window and does not disturb the hen or goose retrievals.

For the Hen–Goose–Rooster triplet, the difficulty is confined to the third window and, within it, to a single line: the green feature shared by all animals. At $\text{rng}(20)$ the standard parameters recall the rooster correctly, with no adjustment needed. At every other seed tested, the green line fails to appear in the rooster's window (1.6–2.0 s) unless W_p is raised: a value of 2.3 recovers it at $\text{rng}(10)$, $\text{rng}(30)$, $\text{rng}(40)$, while $\text{rng}(50)$ and $\text{rng}(80)$ need 2.5, and $\text{rng}(70)$ and $\text{rng}(60)$ needs 2.7. The hen and goose windows were unaffected in every case, so the adjustment is local to the rooster's window and does not disturb the rest of the triplet.




Figure 3.7. (A) The plot shows the suppressed shared features of animals, represented by dashed-dotted green lines, at $W_p=2.1$ (standard). (B) However, after slightly raising W_p to 2.3, the shared featured of animals are restored. The behaviour is typical in long range of rng_s except $\text{rng}(20)$ in Hen-Goose-Rooster triplet.

To test whether this difficulty belongs to the rooster itself or to its position in the sequence, the hen and rooster were swapped, placing the rooster first (0.0–0.4 s) and the hen third (1.6–2.0 s), with the goose left in the middle window. With the rooster cued first, its shared features

were evoked correctly at every seed, with no adjustment to W_p needed. The hen, now in the third window, showed the same pattern the rooster had shown before the swap: its green line recalled correctly at every seed except rng(70), where a slightly raised $W_p=2.3$ was again needed to recover it.




Figure 3.8 Rooster-Goose-hen at rng(10) with standard parameters

This points to the position in the sequence, not the identity of the animal, as the source of the difficulty. Whichever animal is cued last has the least time remaining for its shared, all-animal feature to settle into synchrony before the spotlight window closes, and it is this timing pressure, rather than anything specific to the rooster's own trained connections, that produces the occasional need for a higher W_p in the final window.

3.3 SUMMARY

The simulations show that the recall network with the spotlight mechanism works across a wide range of concepts and seeds. A single cue is enough to reconstruct an attended concept, and the spotlight moves from one concept to the next as the schedule advances and the active representation fatigues, giving a clean sequence of separate retrievals with no carry-over between windows. The mechanism is not tied to any particular set of animals: it recalls arbitrary triplets, and for most concepts the result holds across all eight seeds, with only small per-seed changes to the recurrent excitation needed to keep each case clean. This manual, per-seed tuning is a limitation of the model rather than something it resolves on its own.

The failures that remain fall almost entirely on the giraffe and the zebra, the two most heavily shared concepts. Completing either one requires strong recurrent excitation, but a well-separated concept, such as the lion saturates when driven by that same excitation, so the two demands cannot both be met with one value of W_p . Raising the recurrent inhibition W_f eases the saturation but suppresses the shared features in return, so it does not resolve the conflict. This excitation-versus-saturation trade-off is the main limit on the model's recall, and it is discussed further in the next chapter.

CHAPTER 4: DISCUSSION

This thesis joins two processes that are usually modelled in isolation: the retrieval of concepts from semantic memory and the selective attention that governs which concept is accessed. The model brings them together in a single neural mass network of three regions: a semantic store, a dedicated alpha oscillator, and an auto-associative recall area, and adds a moving attentional spotlight and a synaptic fatigue mechanism. The simulations showed that a single cued feature is enough to recall a whole concept. Attention moves from one concept to the next, and the sequence stays clean, with no earlier retrieval bleeding into a later one. This chapter interprets those results, states the limitations, and outlines where the model could go next.

4.1 APPLYING ALPHA GATING TO A TRAINED SEMANTIC MEMORY

The central result is that the alpha-gating mechanism operates on a trained semantic memory. In (Ursino, 2025) the alpha gate was demonstrated on sensory units and on a regular spatial lattice, where the units carried no learned content, and the task was the suppression of an externally imposed stimulus. Here, the same phase-opposition mechanism acts on 102 units whose connections were learned by the asymmetric Hebb rule of (Ursino & Pirazzini, 2024) so that releasing the gate on a single feature drives the recurrent network to reconstruct the entire concept that feature belongs to. The selection and the completion are therefore handled by different parts of the model working together: the spotlight chooses which stored concept reaches the recall area, and the trained recurrence in that area fills in the rest. Because the gate can be released on any feature, and so on any concept, the selection is general rather than tied to a particular stimulus, which is the property a model of selective semantic access requires.

The result also confirms, in a larger and more structured network, the phase-opposition principle the gating rests on. Suppression depends on the alpha reaching the store roughly 180 degrees out of phase with the alpha reaching the recall area, and the model achieves this through the transmission delay between the two regions. The mechanism behaves as (Ursino, 2025) predicted, and it continues to do so when the suppressed units carry learned semantic content rather than a single imposed stimulus. This integration of a trained semantic network with alpha-based attentional control was identified as future work by both (Ursino & Pirazzini, 2024) and (Ursino, 2025), and the present model implements it.

4.2 SEQUENTIAL RETRIEVAL AND THE ROLE OF FATIGUE

Alpha gating on its own selects a concept, but it does not release it. A strongly activated attractor tends to persist, and without some way of letting it go, the next retrieval would compete with the residue of the last. The synaptic fatigue added to the recall area supplies that release. As a concept's units stay active, their depletion variables fall, the recurrent excitation that holds the concept together weakens, and the representation fades on its own. By the time the spotlight moves to the next cue, the previous concept has lost the recurrent support that

would let it reactivate, and the result is a clean sequence of separate retrievals from minimal input. The fatigue operates on a timescale of a few hundred milliseconds, much slower than the gamma oscillations that carry the features or the alpha rhythm that gates them, so it shapes the turnover between concepts without disturbing the dynamics within a single retrieval. The use of short-term synaptic depression to destabilize an active attractor and move a network follows the same principle used by (Singh, 2022) to drive transitions during memory consolidation, applied here to attentional turnover rather than consolidation.

4.3 ROBUSTNESS AND THE STRUCTURE OF THE DIFFICULT CASES

Recall was robust to the random seed. Across eight seeds, a single setting of the recurrent excitation recovered the large majority of cases, and the few exceptions were corrected by a small change to the excitation or the noise level rather than by any change to the mechanism, which indicates that the behaviour is a property of the model and not of a particular noise realization. The cases that did fail are themselves informative. The hardest grouping combined concepts that place opposing demands on the recurrent network: animals such as the giraffe and zebra carry large and entangled sets of shared-feature memberships and need strong recurrent excitation to complete, while a well-separated concept driven by that same strong excitation is pushed toward saturation. No single excitation value satisfies both at once. This tension is a direct expression of the excitation-inhibition balance identified as central to gamma-band recall by (Ursino et al., 2010) and (Ursino & Pirazzini, 2024), which proved useful only as a targeted remedy for saturation and, at full strength, suppressed features that should have been recalled. The reappearance of the giraffe's "it has spots" feature as a hard case, already reported by (Ursino & Pirazzini, 2024), shows that the difficulty follows the structure of the stored concepts and is inherited from the trained network rather than introduced by the attention mechanism.

4.4 LIMITATIONS

Several limitations qualify these conclusions. The semantic store is a single-region network of four population units without the six-layer laminar structure of more detailed cortical models, so it captures the dynamics of an area rather than its internal organization. The recurrent excitation required minor per-seed tuning; although a single scaling value recalled the great majority of concepts correctly, a few seeds needed a small adjustment to that value to keep the retrieval clean. Features are supplied directly as input, which means the model presumes an earlier processing stage that extracts those features from sensory data and does not itself account for how or where that extraction takes place. The synaptic fatigue is governed by a single global recovery time constant rather than by depletion dynamics derived from physiological measurements. Finally, the balance between recurrent excitation and inhibition was characterized but not resolved: the results establish that inhibition must be set against

excitation to control saturation, yet they do not yield a principled rule for choosing its strength, which at present remains a matter of manual adjustment.

4.5 FUTURE WORK

These limitations point to several directions for future work.

- 1) The recurrent excitation scaling, W_p , which here required manual per-seed adjustment, could instead be set during training, removing the need for manual tuning afterward.
- 2) The neocortex has a six-layer structure, and these layers play different roles in sending and receiving information (Thomson & Bannister, 2003). The present model does not attempt this level of detail. It uses four populations with no layer structure at all: the pyramidal population's output simply feeds into the pyramidal and fast GABAergic populations of another unit, based on what the network learned during training. Building a version with the full six-layer arrangement is left for future work.
- 3) The present model uses just four neuron populations to represent a single oscillating unit, minimal circuitry, built to demonstrate that the proposed mechanism works, not to capture every detail of cortical structure. More recent models (Bensaid, 2019; Wendling, 2024) go further, drawing on current neurophysiological findings to add more types of interneurons, capture microcircuits at the cellular scale, and represent communication across cortical layers, closer to the neocortex's actual layered structure. These give a more faithful account of neuroelectrical activity.
- 4) Attention in the present model works as a ON/OFF switch: a feature's connection to the alpha gate is set to either 300, fully suppressed, or 0, fully attended, with nothing in between. Real attention is unlikely to work this way; alpha amplitude is thought to rise and fall more gradually, allowing partial or shifting attention rather than an all-or-nothing state. Replacing the on/off gate with a continuously graded amplitude would make the attentional dynamics more realistic.
- 5) The fatigue mechanism currently uses a single recovery time constant, $\tau_d=1$ s, chosen because it produced good retrieval in simulation rather than because it matches any measured biological value. Real synapses show short-term depression with recovery rates that have actually been recorded experimentally, how quickly the pool of available neurotransmitter depletes under repeated firing and how quickly it refills at rest. Grounding the model's depletion and recovery rates in these measurements would tie the fatigue mechanism to real synaptic physiology.

- 6) In the present model, the recurrent inhibition W_f is a single fixed value applied equally across all units, and the Giraffe–Zebra–Lion case showed the limits of this approach: raising W_f enough to stop the lion from saturating also suppressed the zebra's shared features, so no one value satisfied both. A more principled approach would make inhibition adaptive, scaling it to the saturation level of each unit rather than applying the same value everywhere, so that units close to saturating receive stronger inhibition while units that are not do not. This would resolve the excitation-inhibition trade-off that the present model can only manage by manually testing different values of W_f and settling for a compromise.
- 7) Reduced fast GABAergic interneuron activity is known to weaken the gamma rhythm and impair the mechanisms that depend on it (Ursino, 2025). Since gating in the present model relies on gamma-band activity being reliably suppressed or released, a similar reduction in fast-inhibitory connectivity would be expected to degrade attentional gating in the same way. This parallels the deficits reported in schizophrenia, and the present three-region architecture offers a natural setting to simulate it, by lowering the fast-inhibitory connectivity and observing how recall and selective attention break down.

4.6 SUMMARY

The work showed that an alpha-based attentional gate can be applied to a trained semantic memory, so that a moving spotlight selects and recalls stored concepts one at a time, and that a synaptic fatigue mechanism lets the network move between them without external intervention. By bringing the storage and the selection of concepts into a single architecture, the model connects semantic retrieval and selective attention in one mechanistic account. Its failures, concentrated in the most heavily shared concepts, mark out an interpretable limit set by the structure of the stored knowledge and the balance between excitation and inhibition.

REFERENCES

- Abbott. (1997). Synaptic depression and cortical gain control. *Science*.
<https://doi.org/10.1126/science.275.5297.220>
- Bacigalupo & Luck. (2022). *Alpha-band EEG suppression as a neural marker of sustained attentional engagement to conditioned threat stimuli*.
- Bensaid. (2019). *A neural mass model for the human neocortex including interneuron-specific circuits*.
- Bonnefond. (2017). Communication between brain areas based on nested oscillations. *eNeuro*.
<https://doi.org/10.1523/ENEURO.0153-16.2017>
- Busch. (2009). The phase of ongoing EEG oscillations predicts visual perception. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.0113-09.2009>
- Buzsaki & Draguhn. (2004). Neuronal oscillations in cortical networks. *Science*.
<https://doi.org/10.1126/science.1099745>
- Chen. (2017). A unified model of human semantic knowledge and its disorders. *Nature Human Behaviour*. <https://doi.org/10.1038/s41562-016-0039>
- David & Friston. (2003). A neural mass model for MEG/EEG: coupling and neuronal dynamics. *NeuroImage*. [https://doi.org/10.1016/S1053-8119\(03\)00264-7](https://doi.org/10.1016/S1053-8119(03)00264-7)
- FitzHugh. (1961). Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*. [https://doi.org/10.1016/S0006-3495\(61\)86902-6](https://doi.org/10.1016/S0006-3495(61)86902-6)
- Fries. (2015). Rhythms for Cognition: Communication through Coherence. *Neuron*.
<https://doi.org/10.1016/j.neuron.2015.09.034>
- Garagnani. (2017). A neuroanatomically grounded Hebbian-learning model of attention-language interactions in the human brain. *Cognitive Computation*.

- Garagnani & Pulvermuller. (2016). Conceptual grounding of language in action and perception: A neurocomputational model of the emergence of category specificity and semantic hubs. *European Journal of Neuroscience*. <https://doi.org/10.1111/ejn.13145>
- Gray. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*. <https://doi.org/10.1038/338334a0>
- Herrmann. (2004). Cognitive functions of gamma-band activity: Memory match and utilization. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2004.06.006>
- Hodgkin & Huxley. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*. <https://doi.org/10.1113/jphysiol.1952.sp004764>
- Ikkai. (2016). Aging and the allocation of attention: Stronger alpha modulation in older adults. *Neurobiology of Aging*.
- Jansen & Rit. (1995). Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns. *Biological Cybernetics*. <https://doi.org/10.1007/BF00199471>
- Janssens. (2018). Occipital alpha power reveals fast attentional inhibition of incongruent distractors. *Psychophysiology*. <https://doi.org/10.1111/psyp.13011>
- Jensen. (2012). An oscillatory mechanism for prioritizing salient unattended stimuli. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2012.03.002>
- Jensen. (2014). Temporal coding organized by coupled alpha and gamma oscillations prioritize visual processing. *Trends in Neurosciences*. <https://doi.org/10.1016/j.tins.2014.04.001>
- Jensen & Colgin. (2007). Cross-frequency coupling between neuronal oscillations. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2007.05.003>
- Kilpatrick, Z. P. (2013). Wilson-Cowan Model. In D. Jaeger & R. Jung (Eds), *Encyclopedia of Computational Neuroscience* (pp. 1–5). Springer New York. https://doi.org/10.1007/978-1-4614-7320-6_80-1

- Klimesch. (2012). Alpha-band oscillations, attention, and controlled access to stored information. *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2012.10.007>
- Lopes da Silva. (1974). Model of brain rhythmic activity: The alpha-rhythm of the thalamus. *Kybernetik*. <https://doi.org/10.1007/BF00271629>
- Mathewson. (2009). To see or not to see: Prestimulus alpha phase predicts visual awareness. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.3963-08.2009>
- McRae. (1997). *Models of semantic memory*.
- Nagumo. (1962). An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*. <https://doi.org/10.1109/JRPROC.1962.288235>
- Nobre & Stokes. (2019). Premembering Experience: A Hierarchy of Time-Scales for Proactive Attention. *Current Opinion in Psychology*. <https://doi.org/10.1016/j.copsyc.2019.04.008>
- Pfurtscheller. (1996). Event-related synchronization (ERS) in the alpha band—An electrophysiological correlate of cortical idling: A review. *International Journal of Psychophysiology*. [https://doi.org/10.1016/S0167-8760\(96\)00066-9](https://doi.org/10.1016/S0167-8760(96)00066-9)
- Quax. (2017). Top-down control of cortical gamma-band communication via pulvinar induced phase shifts in the alpha rhythm. *PLoS Computational Biology*. <https://doi.org/10.1371/journal.pcbi.1005519>
- Rogers & McClelland. (2008). Précis of Semantic Cognition: A Parallel Distributed Processing Approach. *Behavioral and Brain Sciences*. <https://doi.org/10.1017/S0140525X08003648>
- Singer & Gray. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*. <https://doi.org/10.1146/annurev.ne.18.030195.002233>
- Singh. (2022). A model of autonomous interactions between hippocampus and neocortex driving sleep-dependent memory consolidation. *Proceedings of the National Academy of Sciences*. <https://doi.org/10.1073/pnas.2123432119>

- Tallon-Baudry & Bertrand. (1999). Oscillatory gamma activity in humans and its role in object representation. *Trends in Cognitive Sciences*. [https://doi.org/10.1016/S1364-6613\(99\)01299-1](https://doi.org/10.1016/S1364-6613(99)01299-1)
- Thomson & Bannister. (2003). Interlaminar connections in the neocortex. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/13.1.5>
- Thut. (2006). Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.0875-06.2006>
- Tsodyks & Markram. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*. <https://doi.org/10.1073/pnas.94.2.719>
- Ursino. (2025). Modeling the Role of the Alpha Rhythm in Attentional Processing during Distractor Suppression. *Journal of Cognitive Neuroscience*. https://doi.org/10.1162/jocn_a_02327
- Ursino, Cona, & Zavaglia. (2010). The generation of rhythms within a cortical region: Analysis of a neural mass model. *NeuroImage*. <https://doi.org/10.1016/j.neuroimage.2009.12.084>
- Ursino & Pirazzini. (2024). Construction of a Hierarchical Organization in Semantic Memory: A Model Based on Neural Masses and Gamma-Band Synchronization. *Cognitive Computation*. <https://doi.org/10.1007/s12559-023-10202-y>
- van Diepen. (2019). The functional role of alpha-band activity in attentional processing: The current zeitgeist and future outlook. *Current Opinion in Psychology*. <https://doi.org/10.1016/j.copsyc.2019.03.015>
- Voytek. (2010). Shifts in gamma phase-amplitude coupling frequency from theta to alpha over posterior cortex during visual tasks. *Frontiers in Human Neuroscience*. <https://doi.org/10.3389/fnhum.2010.00191>

- Wendling. (2002). Epileptic fast activity can be explained by a model of impaired GABAergic dendritic inhibition. *European Journal of Neuroscience*. <https://doi.org/10.1046/j.1460-9568.2002.01985.x>
- Wendling. (2024). *Multiscale neuro-inspired models for interpretation of EEG signals in patients with epilepsy*.
- Wildegger. (2017). Preparatory alpha-band oscillations reflect spatial gating independently of predictions regarding target identity. *Journal of Neurophysiology*. <https://doi.org/10.1152/jn.00856.2016>
- Wilson & Cowan. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*. [https://doi.org/10.1016/S0006-3495\(72\)86068-5](https://doi.org/10.1016/S0006-3495(72)86068-5)
- Worden. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. *Journal of Neuroscience*.
- Wostmann. (2019). Alpha oscillations in the human brain implement distractor suppression independent of target selection. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.1954-18.2019>